arXiv:1603.07427v1 [stat.ME] 24 Mar 2016

Penalized Weighted Least Squares for Outlier Detection and Robust Regression

Xiaoli Gao*
Department of Mathematics and Statistics
University of North Carolina at Greensboro
and
Yixin Fang
Department of Population Health
New York University School of Medicine

Abstract

To conduct regression analysis for data contaminated with outliers, many approaches have been proposed for simultaneous outlier detection and robust regression, so is the approach proposed in this manuscript. This new approach is called "penalized weighted least squares" (PWLS). By assigning each observation an individual weight and incorporating a lasso-type penalty on the log-transformation of the weight vector, the PWLS is able to perform outlier detection and robust regression simultaneously. A Bayesian point-of-view of the PWLS is provided, and it is showed that the PWLS can be seen as an example of M-estimation. Two methods are developed for selecting the tuning parameter in the PWLS. The performance of the PWLS is demonstrated via simulations and real applications.

Keywords: Adaptive lasso; M-estimation; Outliers; Stability; Tuning

*The authors gratefully acknowledge *Simons Foundation (#359337, Xiaoli Gao) and UNC Greensboro New Faculty Grant (Xiaoli Gao)*

1 Introduction

In statistics, an outlier is an observation that does not follow the model of the majority of the data. Some outliers may be due to intrinsic variability of the data; this type of outliers should be examined carefully using some subgroup analysis. Other outliers may indicate errors such as experimental error and data entry error; this type of outliers should be down-weighted.

To conduct regression analysis for data contaminated with outliers, one can detect outliers first and then run ordinary regression analysis using the data with the detected outliers deleted (Weisberg 2005), or run some version of robust regression analysis which is insensitive to the outliers (Huber 1973). Alternatively, many approaches have been proposed to simultaneously perform outlier detection and robust regression. See for example, the least median of squares (Siegel 1982), the least trimmed squares (Rousseeuw 1984), S-estimates (Rousseeuw & Yohai 1984), Generalized S-estimates (Croux et al. 1994), MM-estimates (Yohai 1987), the robust and efficient weighted least squares estimators (Gervini & Yohai 2002), and forward search (Atkinson et al. 2003). Boente et al. (2002) also studied outlier detection under principal components model. One can refer to Maronna et al. (2006) and Hadi et al. (2009) for broader reviews of some recent robust regression procedures and outlier detection procedures.

In this manuscript, we propose a new approach, penalized weighted least squares (PWLS). By assigning each observation an individual weight and incorporating a lasso-type penalty on the log-transformation of the weight vector, the PWLS is able to perform outlier detection and robust regression simultaneously. For this aim, assume the data are from the following model,

$$y_i = \mathbf{x}_i'\boldsymbol{\beta}^* + \varepsilon_i/w_i^*, \tag{1}$$

where $\mathbf{x}_i \in \mathbb{R}^p$ are the predictor vectors, $\boldsymbol{\beta}^* = (\beta_1^*, \cdots, \beta_p^*)'$ is the coefficient vector, and ε_i with $E(\varepsilon_i|\mathbf{x}_i) = 0$ are random errors following some unknown distribution $F(\cdot|\mathbf{x}_i)$ independently. Also assume the data are contaminated with outliers, and therefore the underlying weights w_i^* are introduced, with $w_i^* < 1$ indicating outliers and $w_i^* = 1$ indicating non-outliers. We shall start our discussion with the homogeneous setting where $Var(\varepsilon_i) = \sigma^2$ for $1 \le i \le n$, and then generalize it to the heterogeneous setting where $\varepsilon_i = g(\mathbf{x}_i'\boldsymbol{\theta})\epsilon_i$ with $E(|\epsilon_i|) = 1$ for $1 \le i \le n$.

In ordinary least squares (OLS) regression, suspected outliers could be visualized by plotting

the residual $r_i = y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}}$ against the predicted outcome $\widehat{y}_i$, where $\widehat{\boldsymbol{\beta}}$ is the estimate of $\boldsymbol{\beta}$, along with other visualizing tools such as studentized-residuals plot and Cook's distance plot (Weisberg 2005). However, when there are multiple outliers, these simple methods can fail, because of two phenomena, masking and swamping. These two phenomena can be demonstrated by examining a famous artificial dataset, the Hawkins-Bradu-Kass (HBK) data (Hawkins et al. 1984).

The HBK data consist of 75 observations, where each observation has one outcome variable and three predictors. The first 14 observations are leverage points; however, only the first 10 observations are actual outliers. The studentized residual plot is shown in the left panel of Figure 1, where those from the 10 actual outliers are displayed differently. Those observations with large residuals bigger than some threshold are suspected to be outliers, and the threshold 2.5 suggested by Rousseuw & Leroy (1987) is also shown in the residual plot. It is shown that three non-outliers (displayed in light dot) are beyond the threshold lines and therefore are suspected to be outliers; this is the swamping phenomenon. It is also shown that seven outliers (displayed in dark asterisk) are within the threshold lines and therefore survive the outlier screening; this is the masking phenomenon.

The other two plots in Figure 1 are the outputs of our new method, the PWLS. As tuning parameter λ goes from zero to infinity, it generates a weight path for each observation. These weight paths are shown in the middle panel; solution paths of non-outliers and outliers are in light solid curves and dark dashed curves, respectively. We can see that paths of the ten actual outliers are distant from the others, and at the selected tuning parameter $\widehat{\lambda}$ (displayed in a vertical line), weights of those non-outliers are exactly equal to one while the weights of those ten outliers are very small. The choice of optimal tuning parameter is to be presented in Subsection 3.2, where a random-weighting procedure is developed to estimate the probability of each observation being outlier at each value of tuning parameter λ. For the HBK data, such probabilities along a wide range of λ are shown in the right panel; light solid curves and dark dashed curves are for non-outliers and outliers, respectively. We can see that the estimated outlier probabilities of those ten outliers are much higher than the others, and at the same $\widehat{\lambda}$ (vertical line), the probabilities from non-outliers are exactly equal to or at least close to 0 while those from ten outliers are close to 1.

The proposal of our new method is motivated by a seminal paper, She & Owen (2011). In their paper, a regression model with a mean shift parameter is considered and then a lasso-type penalty

Figure 1: An illustrative example using HBK data. Left: Studentized residual plot from LS regression, with threshold lines of ± 2.5 (non-outliers displayed as light dots and outliers displayed as dark asterisks); Middle: PLWS weight solution paths with a selected tuning parameter shown in a vertical line (non-outliers displayed as light solid lines and outliers displayed as dark dashed lines); Right: Outlier probability paths generated using random weighting with a selected tuning parameter shown in a vertical line (non-outliers displayed as light solid lines and outliers displayed as dard dashed lines).

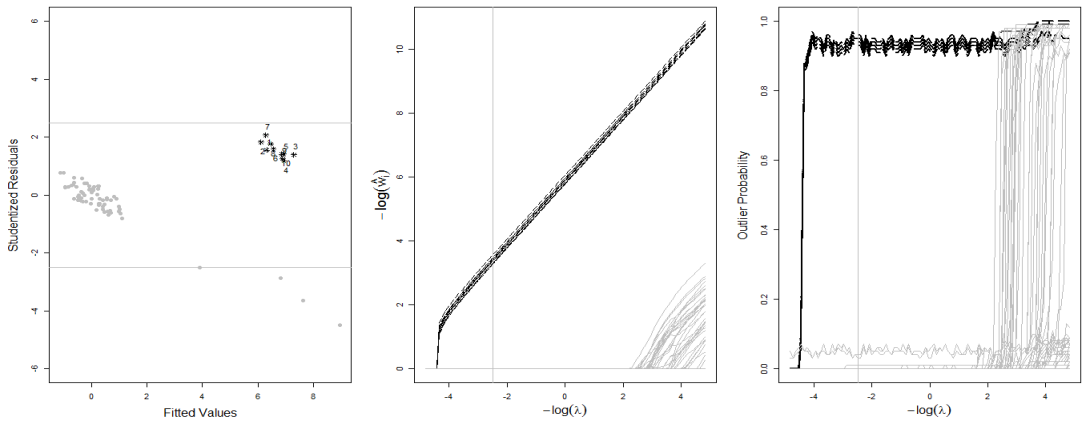

is incorporated. However, the soft-thresholding implied by the lasso-type penalty cannot counter the masking and swamping effects and therefore they introduced a hard-thresholding version of their method. Surprisingly, this small change from soft-thresholding to hard-thresholding made their method work well for countering the masking and swamping effects, although the mysterious reason behind this was not uncovered.

The remaining manuscript is organized as follows. In Section 2, we discuss the PWLS, along with some of model justification including its Bayesian understanding and robust investigation. In Section 3, we develop an algorithm to implement the proposed method, and two methods for selecting the tuning parameter in it. In Section 4, we extend the PWLS to heterogeneous models, in particular, the variance function linear models. In Section 5, we evaluate the performance of the newly proposed method using simulations and real applications. Some discussion is in Section 6 and the technical proof is relegated to the Appendix.

2 The Penalized Weighted Least Squares

If the weight vector $\mathbf{w}^* = (w_1^*, \cdots, w_n^*)'$ in model (1) is given in advance, $\boldsymbol{\beta}^*$ can be estimated by minimizing the weighted sum of squares, $\sum_{i=1}^n w_i^{*2}(y_i - \mathbf{x}_i'\boldsymbol{\beta})^2$. In the approach we develop here, we allow weights to be data-driven and estimate both $\boldsymbol{\beta}^*$ and $\mathbf{w}^*$ jointly by minimizing the following penalized weighted least squares (PWLS),

$$(\widehat{\boldsymbol{\beta}}, \widehat{\mathbf{w}}) = \arg\min_{\boldsymbol{\beta}, \mathbf{w}} \left\{ \sum_{i=1}^n w_i^2 (y_i - \mathbf{x}_i'\boldsymbol{\beta})^2 + \sum_{i=1}^n \lambda |\log(w_i)| \right\}, \tag{2}$$

where tuning parameter λ controls the number of suspected outliers. The non-differentiability of penalty $|\log(w_i)|$ over $w_i = 1$ implies that some of the components of $\widehat{\mathbf{w}}$ may be exactly equal to one. Then the observations corresponding to $\widehat{w}_i = 1$ survive the outlier screening, while those corresponding to $\widehat{w}_i \neq 1$ are suspected to be outliers. Therefore, the PWLS can perform simultaneous outlier detection and robust estimation.

Noting that $|\log(w_i)| = |\log(1/w_i)|$, we can assume that all the components of $\mathbf{w}$ are either less than one (suspected outliers) or equal to one (non-outliers). In fact, any $w_i > 1$ must not be a solution since it can be always replaced by $\widetilde{w}_i = 1/w_i < 1$ and decreases the objective function. Therefore, in the first term of the objective function of (2), those suspected outliers are assigned smaller weights than the others.

The performance of the PWLS depends crucially on the determination of tuning parameter λ, ranging from 0 to ∞. When λ is sufficiently large, all $\log(\widehat{w}_i)$ become zero, and consequently all observations survive outlier screening. When λ is sufficiently small, some $\widehat{w}_i$ become zero, and consequently they could be suspected as outliers. Therefore, the tuning parameter selection plays an important role in determining the amount of outliers. Two methods for tuning parameter selection are discussed in the next section.

Finally, we should emphasize that the PWLS is not a variation of the classical weighted least squares (WLS; see e.g., Weisberg 2005) aiming for fitting heterogeneous data. The PWLS is coined because the first term in the objective function of (2), the weighted sum of squares, is the same as that for the WLS. We could conduct variable selection by adding some penalty term on the regression coefficients in the WLS and also call it the penalized weighted least squares, but the readers should not be confused by these names, keeping in mind that the goal of this

manuscript is outlier detection rather than variable selection.

2.1 A Bayesian understanding of the PWLS

We provide a Bayesian understanding of model (2). Denote $\nu_i = 1/w_i$ and $\boldsymbol{\nu} = (\nu_1, \cdots, \nu_n)'$. Let $\pi(\boldsymbol{\beta})$, $\pi(\sigma^2)$, and $\pi(\nu_i)$ be the independent prior distributions of $\boldsymbol{\beta}$, σ^2, and ν_i, respectively. Assume non-informative priors $\pi(\boldsymbol{\beta}) \propto 1$ and $\pi(\sigma^2) \propto 1/\sigma^2$ for $\boldsymbol{\beta}$ and σ^2, respectively. Also assume that ν_i has a Type I Pareto distribution with hyper-parameter $\lambda_0 \geq 1$; that is,

$$\pi(\nu_i) \propto \nu_i^{1-\lambda_0} I(\nu_i \geq 1), \quad \text{for } 1 \leq i \leq n, \tag{3}$$

where $I(\cdot)$ is the indicator function. The prior distribution of ν_i with different hyper-parameters is shown in the left panel of Figure 2. In particular, it is the uniform non-informative prior when $\lambda_0 = 1$ and Jeffreys non-informative prior when $\lambda_0 = 2$.

Then the joint posterior distribution of the parameters is equal to

$$\pi(\boldsymbol{\beta}, \sigma^2, \boldsymbol{\nu}|\mathbf{y}) \propto (\sigma^2)^{-n/2-1} \prod_{i=1}^{n} \nu_i^{-\lambda_0} \exp\left\{ -\frac{1}{2\sigma^2} \sum_{i=1}^{n} \frac{1}{\nu_i^2} (y_i - \mathbf{x}_i'\boldsymbol{\beta})^2 \right\}.$$

The mode, $(\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\nu}})$, of the above posterior distribution is

$$(\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\nu}}) = \arg\min_{\boldsymbol{\beta}, \boldsymbol{\nu}} \left\{ \sum_{i=1}^{n} \frac{1}{\nu_i^2} (y_i - \mathbf{x}_i'\boldsymbol{\beta})^2 + \sum_{i=1}^{n} 2\widehat{\sigma}^2 \lambda_0 |\log(\nu_i)| \right\}, \tag{4}$$

where $\widehat{\sigma}^2 = (n+2)^{-1} \sum_{i=1}^{n} \widehat{\nu}_i^{-2} (y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}})^2$. Thus (4) is equivalent to (2) if $\lambda = 2\widehat{\sigma}^2\lambda_0$.

2.2 In connection with M-estimation

We demonstrate that the PWLS is an example of M-estimation by deriving its implicit ψ and ρ functions. Consider M-estimation with ψ function,

$$\psi(t, \lambda) = \begin{cases} \lambda/t, & \text{if } |t| > \sqrt{\lambda/2}, \\ 2t, & \text{if } |t| \leq \sqrt{\lambda/2}, \end{cases} \tag{5}$$

Figure 2: Display of some functions. Left: Improper priors of v_i with hyper-parameter $\lambda_0 = 1, 2, 3, \infty$; Middle: The ρ function with tuning parameter $\lambda = 1, 2, 3$. Right: The ψ function with tuning parameter $\lambda = 1, 2, 3$.

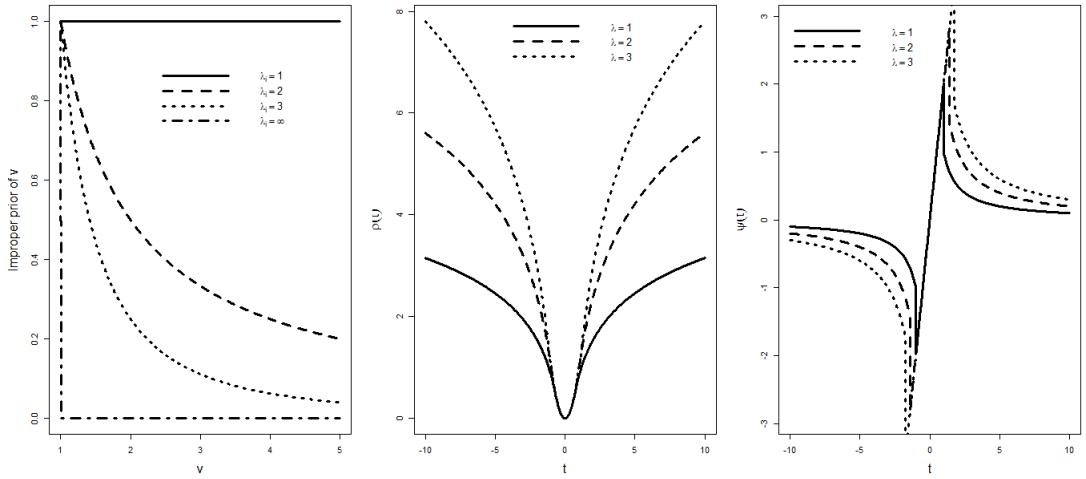

and the corresponding ρ function,

$$\rho(t, \lambda) = \begin{cases} \lambda \log(|t|\sqrt{2/\lambda}) + \lambda/2, & \text{if } |t| > \sqrt{\lambda/2}, \\ t^2, & \text{if } |t| \leq \sqrt{\lambda/2}. \end{cases} \tag{6}$$

The above ρ and ψ functions with different λ are displayed in the middle panel and right panel of Figure 2, respectively. Apparently, the proposed ψ function in (5) is non-decreasing near the origin, but decreasing toward 0 far from the origin. Therefore, (6) generates a redescending M-estimator having special robustness properties.

Consider the M-estimator with a concomitant scale, where the concomitant scale is added to ensure that we can estimate $\boldsymbol{\beta}$ and σ simultaneously,

$$(\widehat{\boldsymbol{\beta}}_M, \widehat{\sigma}_M) = \underset{\boldsymbol{\beta}, \sigma}{\arg\min} \left\{ \sum_{i=1}^{n} \rho\left(\frac{y_i - \mathbf{x}_i' \boldsymbol{\beta}}{\sigma}, \lambda \right) + 2cn \log \sigma \right\}. \tag{7}$$

With the proof in Appendix, we show that, for any λ and c, estimator $\widehat{\boldsymbol{\beta}}_M$ from (7) is the same

as that from the following PWLS with the same concomitant scale c,

$$(\widehat{\boldsymbol{\beta}}_P, \widehat{\sigma}_P, \widehat{\mathbf{w}}_P) = \underset{\boldsymbol{\beta},\sigma,\mathbf{w}}{\arg\min} \left\{ \frac{1}{\sigma^2} \left[\sum_{i=1}^{n} w_i^2 (y_i - \mathbf{x}_i'\boldsymbol{\beta})^2 + \sum_{i=1}^{n} \lambda |\log(w_i)| \right] + 2cn \log \sigma \right\}. \tag{8}$$

Theorem 1 *For any given λ and c, the M-estimator $\widehat{\boldsymbol{\beta}}_M$ from* (7) *is the same as the PWLS estimator $\widehat{\boldsymbol{\beta}}_P$ from* (8).

2.3 The adaptive PWLS

Since the lasso was proposed by Tibshirani (1996), many other sparsity inducing penalties have also been proposed. Among them, the adaptive lasso proposed by Zou (2006) has become very popular recently, partly because of its convexity, selection consistency and oracle property. If we see the penalty in (2) as a lasso-type penalty, then we propose the following adaptive version of the PWLS (aPWLS),

$$(\widehat{\boldsymbol{\beta}}, \widehat{\mathbf{w}}) = \underset{\boldsymbol{\beta},\mathbf{w}}{\arg\min} \left\{ \sum_{i=1}^{n} w_i^2 (y_i - \mathbf{x}_i'\boldsymbol{\beta})^2 + \sum_{i=1}^{n} \lambda \varpi_i |\log(w_i)| \right\}, \tag{9}$$

where λ is a tuning parameter controlling the number of suspected outliers and $\boldsymbol{\varpi} = (\varpi_1, \cdots, \varpi_n)'$ includes pre-defined penalty scale factors for all observations. In particular, we expect ϖ_i to be larger for potential outliers and smaller for normal observations. For example, we can take $\varpi_i = 1/|\log(w_i^{(0)})|$, where $w_i^{(0)}$ are some initial estimates of w_i.

The selection of initial estimates $w_i^{(0)}$ is important and therefore we try to make the selection process as objective as possible. First, we obtain $w_i^{(0)}$ using (2) with $\lambda^{(0)} = 2(\widehat{\sigma}^{(0)})^2$. This tuning parameter is suggested in Subsection 2.1 assuming the uniform non-informative prior of (3). Then we propose to consider $\varpi_i = 1/|\log(w_i^{(0)})|$, with the convention that $1/0$ equals some large number, say 999. Based on our limited numerical experience, we find that the performance of the PWLS is robust to a wide range of $\lambda^{(0)}$, as long as the proportion of $w_i^{(0)} = 1$ in the resulting $\mathbf{w}^{(0)}$ is not very high (i.e., as long as it is smaller than 1 minus the proportion of "suspected" outliers).

3 Implementation and Tuning

3.1 Algorithms

We describe an algorithm to implement the aPWLS, of which the PWLS is a specification with $\varpi_i = 1$. Note that the objective function in (9) is bi-convex; For a given $\mathbf{w}$, the function of $\boldsymbol{\beta}$ is a convex optimization problem, and the vice versa. This biconvexity guarantees that the algorithm has promising convergence properties (Gorski et al. 2007). The algorithm is summarized in the following flow-chart.

Algorithm 1 The PWLS

Given $\mathbf{X} \in \mathbb{R}^{n\times p}$, $\mathbf{y} \in \mathbb{R}^n$, initial estimates $\boldsymbol{\beta}^{(0)}$, $\mathbf{w}^{(0)}$ and penalty scales ϖ.

For any given λ in a fine grid, let $j = 1$.

While not converged **do**

[update $\boldsymbol{\beta}$]

$\mathbf{y}^{\text{adj}} = \mathbf{w}^{(j-1)} \cdot \mathbf{y}$, $\mathbf{X}^{\text{adj}} = \mathbf{w}^{(j-1)} \cdot \mathbf{X}$, $\boldsymbol{\beta}^{(j)} = (\mathbf{X}^{\text{adj}'}\mathbf{X}^{\text{adj}})^{-1}\mathbf{X}^{\text{adj}'}\mathbf{y}^{\text{adj}}$

[update $\mathbf{w}$]

$\mathbf{r}^{(j)} = \mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(j)}$,

If $|r_i^{(j)}| > \sqrt{0.5\lambda_i}$, let $\mathbf{w}_i^{(j)} \leftarrow \sqrt{0.5\lambda_i}/|r_i^{(j-1)}|$

otherwise $\mathbf{w}_i^{(j)} \leftarrow 1$

converged $\leftarrow \|\mathbf{w}^{(j)} - \mathbf{w}^{(j-1)}\|_\infty < \epsilon$

$j \leftarrow j + 1$

end while

deliver $\widehat{\boldsymbol{\beta}} = \boldsymbol{\beta}^{(j)}$ and $\widehat{\mathbf{w}} = \mathbf{w}^{(j)}$.

In addition, the corresponding R codes are available at `https://sites.google.com/a/uncg.edu/xiaoli-gao/home/r-code`. The algorithm for (9) is illustrated in the middle panel of Figure 1 using the HBK data, where the paths of $\widehat{\mathbf{w}}$ as λ changes are displayed.

3.2 Tuning parameter selection

We propose two methods for selecting the tuning parameter in the aPWLS; one is Bayesian Information Criterion (BIC; Schwarz (1978)) and the other is based on stability selection (Sun

et al. 2013). Both methods are used to select an "optimal" $\widehat{\lambda}$ from a fine grid of λ.

Let $\widehat{\boldsymbol{\beta}}(\lambda)$ and $\widehat{\mathbf{w}}(\lambda)$ be the resulting estimates for given λ. The BIC method chooses the optimal $\widehat{\lambda}$ that minimizes

$$\mathrm{BIC}(\lambda) = (n-p)\log\left\{\|\widehat{\mathbf{w}}(\lambda)\cdot(\mathbf{y}-\mathbf{X}\widehat{\boldsymbol{\beta}}(\lambda))\|_2^2/\|\widehat{\mathbf{w}}(\lambda)\|_2^2\right\} + \widehat{k}(\lambda)\{\log(n-p)+1\}, \quad (10)$$

where "$\cdot$" is a dotted product and $\widehat{k}(\lambda) = \#\{1 \le i \le n : \widehat{w}_i(\lambda) < 1\}$. The first term in (10) measures the goodness of fit, and the second term measures the model complexity, where $\widehat{k}(\lambda)$ is the number of "outliers" detected using the current tuning parameter λ. The BIC formula indicates a trade-off between the goodness of fit and the number of suspected outliers, with smaller λ leading to more suspected outliers and vice versa. A very similar formula of BIC was also used by She & Owen (2011) for the tuning parameter selection in their methods.

The stability selection method is motivated by a notion that an appropriate tuning parameter should lead to stable outputs if the data are perturbed. In the aPWLS, one of the main outputs is which of n observations are suspected outliers. That is, given λ, inputting data $\mathcal{Z}$ outputs a subset, $\mathcal{O}(\lambda;\mathcal{Z})$, consisting of all the suspected outliers. If there are two perturbed datasets, $\mathcal{Z}^{*1}$ and $\mathcal{Z}^{*2}$, we hope that the two outputs, $\mathcal{O}(\lambda;\mathcal{Z}^{*1})$ and $\mathcal{O}(\lambda;\mathcal{Z}^{*2})$, be similar. Otherwise, if $\mathcal{O}(\lambda;\mathcal{Z}^{*1})$ and $\mathcal{O}(\lambda;\mathcal{Z}^{*2})$ are very different, then neither of the two outputs is trustful. Therefore, we attempt to select a λ such that the resulting output $\mathcal{O}(\lambda;\mathcal{Z})$ is most stable when $\mathcal{Z}$ is perturbed.

Before describing the stability selection method, we should first decide which perturbation procedure is appropriate for our setting. There are three popular perturbation procedures (Shao & Tu 1995): data-splitting, bootstrap, and random weighing. Both data-splitting and bootstrap have been widely used for constructing stability selection methods. For example, Meinshausen & Buhlmann (2010) and Sun et al. (2013) used data-splitting for their proposals of stability selection, while Bach (2004) used bootstrap for his proposal of stability selection. However, neither data-splitting nor bootstrap is suitable for our purpose of outlier detection, because any perturbed dataset using either data-splitting or bootstrap leaves out some observations, whose statuses of being suspected outliers are unobtainable. Therefore, we propose to use random weighting as the perturbation procedure in the construction of our stability selection method. Here the random weighting method, which is a resampling method acting like the bootstrap, is not a Bayesian method, although it was called the Bayesian bootstrap in Rubin (1981).

Let $\omega_1, \cdots \omega_n$ be some i.i.d. random weights with $E(\omega_i) = Var(\omega_i) = 1$, and $\boldsymbol{\omega} = (\omega_1, \cdots \omega_n)'$. Those moment conditions on the random weights are standard (Fang & Zhao 2006). With these random weights, we obtain the corresponding perturbed estimates,

$$(\widehat{\boldsymbol{\beta}}(\lambda; \boldsymbol{\omega}), \widehat{\mathbf{w}}(\lambda; \boldsymbol{\omega})) = \underset{\boldsymbol{\beta}, \mathbf{w}}{\arg\min} \left\{ \sum_{i=1}^{n} \omega_i w_i^2 (y_i - \mathbf{x}_i' \boldsymbol{\beta})^2 + \sum_{i=1}^{n} \lambda \varpi_i |\log(w_i)| \right\}. \tag{11}$$

Via (11), any two sets of random weights, $\boldsymbol{\omega}_1$ and $\boldsymbol{\omega}_2$, give two perturbed weight estimates $\widehat{\mathbf{w}}(\lambda; \boldsymbol{\omega}_1)$ and $\widehat{\mathbf{w}}(\lambda; \boldsymbol{\omega}_2)$, which claim two sets of suspected outliers, $\mathcal{O}(\lambda; \boldsymbol{\omega}_1)$ and $\mathcal{O}(\lambda; \boldsymbol{\omega}_2)$. The agreement of these two sets of suspected outliers can be measured by Cohen's kappa coefficient (Cohen 1960), $\kappa(\mathcal{O}(\lambda; \boldsymbol{\omega}_1), \mathcal{O}(\lambda; \boldsymbol{\omega}_2))$.

Finally, if we repeatedly generate B pairs of random weights, $\boldsymbol{\omega}_{b1}$ and $\boldsymbol{\omega}_{b2}$, $b = 1, \cdots, B$, we can estimate the stability of the outlier detection by

$$\widehat{S}(\lambda) = \frac{1}{B} \sum_{b=1}^{B} \kappa \left(\mathcal{O}(\lambda; \boldsymbol{\omega}_{b1}), \mathcal{O}(\lambda; \boldsymbol{\omega}_{b2}) \right), \tag{12}$$

and then select $\widehat{\lambda}$ that maximizes $\widehat{S}(\lambda)$. As a byproduct and without extra computing, the proposed stability selection method can provide, for each observation, an estimate for the probability of it being an outlier as λ changes,

$$\widehat{P}_i^o(\lambda) = \frac{1}{2B} \sum_{b=1}^{B} \sum_{k=1}^{2} I \left\{ i \in \mathcal{O}(\lambda; \boldsymbol{\omega}_{bk}) \right\}. \tag{13}$$

The stability selection method is illustrated in Figure 1 using the HBK data. The vertical lines shown in the middle and right panels of Figure 1 are corresponding to $\widehat{\lambda}$ selected by the stability selection method. The right panel of Figure 1 shows the outlier probability curves using (13), where the curves of those ten outliers can be distinguished easily from the others.

4 Extension of the PWLS to Heterogeneous Models

Hitherto we consider ε_i in model (1) to be homogeneous and propose the PWLS approach for simultaneously conducting robust regression and detecting outliers. However, when ε_i are also

heterogeneous, and if the heterogeneity is not taken into account, some non-outliers with large underlying variances might be suspected falsely as outliers (the swamping phenomenon), while some outliers with small underlying variances might survive the outlier screening (the masking phenomenon). Therefore, we extend our proposal to be applicable to heterogeneous models.

Consider a heterogeneous case where $\varepsilon_i = g(\mathbf{x}_i'\vartheta)\epsilon_i$ with $E(\epsilon_i) = 0$ and $E(|\epsilon_i|) = 1$ for $1 \le i \le n$. Here we assume $g(v)$ is a known function; for example, $g(v) = |v|$ or $g(v) = \exp(v)$. This is a broad class of heterogeneous models considered in a seminal paper, Davidian & Carroll (1987), where they proposed a general framework for estimating parameter ϑ in variance function $g(\mathbf{x}_i'\vartheta)$. We refer to this class of models as variance function linear models (VFLMs).

Motivated by Davidian & Carroll (1987), many authors have attempted to broaden the class of VFLMs; to name just a few, Carroll & Ruppert (1988), Hall & Carroll (1989), Carroll & Härdle (1989), Carroll (2003), Ma et al. (2006), and Ma & Zhu (2012). Most recently, Lian et al. (2014) studied the variance function partially linear single index models (VFPLSIMs), in which variance function is a function of the sum of linear and single index functions; that is $g(\mathbf{x}_i'\vartheta + h(\mathbf{x}_i'\boldsymbol{\zeta}))$, where g is known and h is unknown. In this manuscript, we demonstrate that we can extend the aPWLS to the VFLMs. Similarly, we can also extend the PWLS to broader and more flexible classes, say the VFPLSIMs, by replacing $g(\mathbf{x}_i'\vartheta)$ in the following discussion by $g(\mathbf{x}_i'\vartheta + h(\mathbf{x}_i'\boldsymbol{\zeta}))$.

Without considering outliers, one of the several approaches proposed in Davidian & Carroll (1987) to estimating $\boldsymbol{\beta}$ and ϑ in the VFLMs is described briefly in three steps. (1) Obtain initial estimate $\widehat{\boldsymbol{\beta}}^{\text{homo}}$ for $\boldsymbol{\beta}$ ignoring heterogeneity; (2) Let $R_i = |y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}}^{\text{homo}}|$ be the absolute residuals and obtain an estimate for ϑ, $\widehat{\boldsymbol{\vartheta}} = \arg\min_{\vartheta} \sum_{i=1}^{n} \left(R_i - g(\mathbf{x}_i'\vartheta)\right)^2$; (3) Obtain an updated estimate for $\boldsymbol{\beta}$, $\widehat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}} \sum_{i=1}^{n} (y_i - \mathbf{x}_i'\boldsymbol{\beta})^2 / g^2(\mathbf{x}_i'\widehat{\boldsymbol{\vartheta}})$. Davidian & Carroll (1987) showed the consistency and efficiency of this method under some regular conditions. Following this approach, we extend the PWLS to the VFLMs for robust regression and outlier detection.

With considering outliers in fitting the VFLM, the extended aPWLS has also three steps. First, ignoring heterogeneity, obtain an initial estimate of $\widehat{\boldsymbol{\beta}}$, $\widehat{\boldsymbol{\beta}}^{\text{homo}}$ using the original aPWLS proposed in Section 2; Second, letting $R_i = |y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}}^{\text{homo}}|$ be the absolute residuals, obtain an estimates for ϑ,

$$\widehat{\boldsymbol{\vartheta}} = \arg\min_{\vartheta} \sum_{i=1}^{n} \left(R_i - g(\mathbf{x}_i'\vartheta)\right)^2; \tag{14}$$

Figure 3: An illustrative example of applying the extended aPWLS to a heterogeneous dataset. Left: Scatter-plot of y_i against $g(\mathbf{x}_i'\boldsymbol{\theta})$ (non-outliers displayed as light circle dots and outliers as dark asterisks). Middle: Plot of Studentized residuals, with threshold lines of ± 2.5 . Right: Weight paths from the extended aPWLS method (non-outliers displayed as light solid lines and outliers as dark dashed lines), with a selected tuning parameter shown in a vertical line.

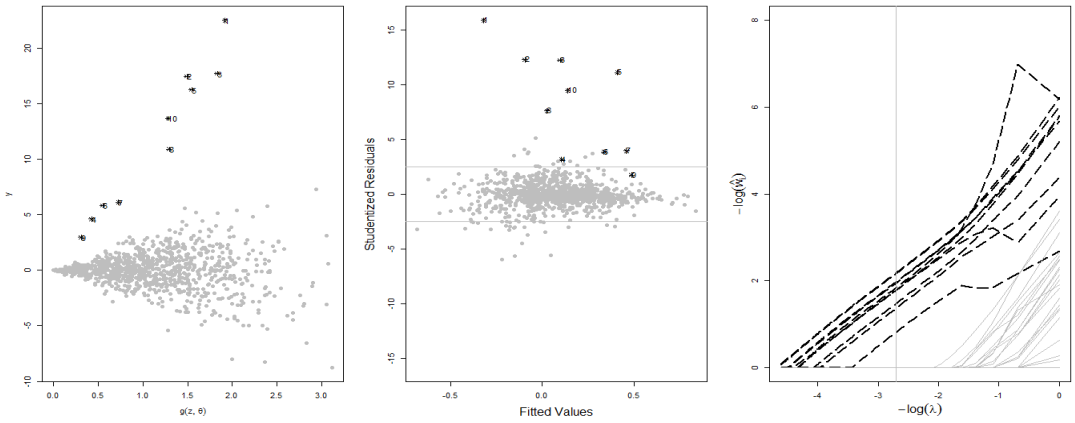

Finally, obtain an updated estimate for $\boldsymbol{\beta}$ and an estimate for $\mathbf{w}$ via,

$$(\widehat{\boldsymbol{\beta}}, \widehat{\mathbf{w}}) = \arg\min_{\boldsymbol{\beta}, \mathbf{w}} \left\{ \sum_{i=1}^{n} w_i^2 (y_i - \mathbf{x}_i'\boldsymbol{\beta})^2 / g^2(\mathbf{x}_i'\widehat{\boldsymbol{\vartheta}}) + \sum_{i=1}^{n} \lambda \varpi_i |\log(w_i)| \right\}. \tag{15}$$

We illustrate this extended aPWLS using a heterogeneous dataset generated from Example 2 described in the next section, where there are 1000 observations and among them 10 observations are outliers. The scatter-plot of y_i vs. $g(\mathbf{x}_i'\vartheta)$ is shown in the left panel of Figure 3, where 10 outliers are displayed in dark asterisks. The studentized residuals are shown in the middle panel of Figure 3. Using threshold 2.5, one outlier is not detected (the masking phenomenon) and many non-outliers are detected falsely as outliers (the swamping phenomenon). The weight paths from the extended aPWLS are shown in the right panel of Figure 3, where the weight paths of 10 outliers (displayed in dark dashed lines) are distinguished from the other from non-outliers (displayed in light solid lines). Moreover, at the selected tuning parameter (displayed in a vertical line), the weights from non-outliers are exactly equal to one while those from outliers are near zero.

5 Numerical Results

As discussed in Section 3.2, tuning parameter selection plays an important role in the penalization approach. We first conduct some simulation studies to compare the two tuning methods, BIC and random weighting (RW). The results show that they perform similarly in terms of outlier detection; the results are omitted here. Such a phenomenon is also observed in Sun et al. (2013). Therefore, because BIC was used with HIPOD in She & Owen (2011), which is the main method with which our method is compared, we use BIC in all the simulation studies presented here. However, RW is used in all three real data applications, because of the byproduct of using the random weighting method, that is, for each observation, we can visualize the probability of it being an outlier as tuning parameter λ changes.

5.1 Simulation studies

We conduct simulation studies to demonstrate the performance of the PWLS for outlier detection under two scenarios, homogeneous models and heterogeneous models. The PWLS is compared with the hard-IPOD (HIPOD) of She & Owen (2011). In She & Owen (2011), the HIPOD was compared with four other robust regression methods. Because Example 1 we consider here is adopted from She & Owen (2011), we are able to compare the PWLS with the HIPOD and at the same time with those four robust regression methods, by combining the results presented here and those presented in She & Owen (2011).

Example 1 (Homogeneous model) *Data are generated from the mean shift model,*

$$\mathbf{y}_i = \mathbf{x}_i'\boldsymbol{\beta} + \gamma_i + \varepsilon_i, \quad 1 \le i \le n, \tag{16}$$

where $\varepsilon_i \sim N(0,1)$ *are independently and* $\boldsymbol{\beta} = \mathbf{1}_p = (1, \cdots, 1)'$. *The first* k *observations are set to be outliers by letting* $\gamma_i = r$ *for* $1 \le i \le k$ *and* 0 *otherwise. A matrix is generated from* $\mathbf{X} = (\mathbf{x}_1, \cdots, \mathbf{x}_n)' = \mathbf{U}\boldsymbol{\Sigma}^{1/2}$, *where* $\mathbf{U} = (u_{ij})_{n\times p}$ *with* $u_{ij} \sim \mathrm{Unif}(-15, 15)$ *and* $\boldsymbol{\Sigma}_{ij} = (\sigma_{ij})_{p\times p}$ *with* $\sigma_{ii} = 1$ *and* $\sigma_{ij} = 0.5$ *for* $i \ne j$. *The design matrix is either* $\mathbf{X}$ *(no* L*) or* $\mathbf{X}$ *with its first* k *rows replaced by* $L \cdot \mathbf{1}_p$ *for some positive* L. *Thus, for the former case (no* L*), the first* k *observations are outliers but not leverage points, whereas for the latter case, the first* k *observations are both leverage points and outliers. Set* $p \in \{15, 50\}$, $k \in \{100, 150, 200\}$, $r = 5$, *and* $L \in \{25, 15, 10\}$.

The tuning parameters in the HIPOD and PWLS are selected via BIC over a grid of 100 values of λ changing from $\lambda_{\max}$ for which no outlier is detected down to $\lambda_{\min}$ at which at least 50% of the observation to be selected as outliers. Here we set $\lambda_{\max} = \|(\mathbf{I}_n - \mathbf{H})\mathbf{y}/\sqrt{\text{diag}(\mathbf{I}_n - \mathbf{H})}\|_\infty$, where $\mathbf{H}$ is the hat matrix and the division is elementwise, and compute solutions along a grid 100 λ values that are equally spaced on the log scale. The initial estimate of $\boldsymbol{\beta}^{(0)}$ is obtained from *lmRob()* in the R package *robust*. The computation of the initial estimate of $w_i^{(0)}$ are described in Section 3.1. In particular, letting $\lambda_0 = \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(0)}\|^2/(n-p)$, then $w_i^{(0)} = \lambda_0/r_i^2$ or 1 if $\lambda_0 < r_i^2$, or otherwise. The simulation results are summarized from 1,000 iterations and they are reported in Table 1.

Similar to She & Owen (2011), we evaluate the outlier detection performance using the mean masking probability (M: fraction of undetected true outliers), the mean swamping probability (S: fraction of non-outliers labeled as outliers), and the joint outlier detection rate (JD: fraction of repetitions with 0 masking) to summarize the results from the 1,000 repetitions. The higher JD is, the better; the smaller M and S are, the better.

From Table 1, we see that the PWLS outperforms the HIPOD in all settings in terms of criteria M and JD; the PWLS has much higher joint outlier detection rate and smaller masking probability. However, the PWLS has a little bit bigger swamping probability measured by S. The comparison is striking when the leverage effect is large ($L = 25$) under 10% outlier ratio ($k = 100$).

Both PWLS and HIPOD loses their efficiency with the existence of a large amount of large leverage points exist ($k = 200$ and $L = 25$). However, PWLS still performs better than the HIPOD especially when the leverage value is is not too big such as ($k = 200$ and $L = 10$).

Example 2 (Heterogeneous model) *Data are generated from the following VFLM,*

$$\mathbf{y}_i = \mathbf{x}_i'\boldsymbol{\beta} + \gamma_i g(\mathbf{z}_i'\vartheta) + \varepsilon_i, 1 \le i \le n, \tag{17}$$

where $\boldsymbol{\beta} = \mathbf{1}_p$, $\mathbf{z}_i = (1, x_{ip})'$, $\vartheta = (1, 0.7)'$ $\varepsilon_i = g(\mathbf{z}_i'\vartheta)\epsilon_i$ *with* $\epsilon_i \sim N(0, 0.5\pi))$. *All* $\mathbf{x}_i$ *are generated independently from a multivariate normal distribution* $N(\mathbf{0}_p, \boldsymbol{\Sigma})$, *where* $\boldsymbol{\Sigma} = (\sigma_{ij})_{p\times p}$ *with* $\sigma_{ij} = 0.5^{|i-j|}$. *The first* k *observations are set bo be outliers by letting* $\gamma_i = r$ *for* $1 \le i \le k$ *and* 0 *otherwise, Set* $n = 1000$, $p \in \{15, 50\}$, $k = 10$, *and* $r = \{20, 15, 10, 5\}$. *We consider the following two heterogeneous settings.*

Table 1: Example 1 − Outlier detection evaluation for homogeneous model (M: the mean masking probability; S: the mean swamping probability; JD: the joint outlier detection rate)

k	p	Method	JD (%)	M (%)	S (%)	JD (%)	M (%)	S (%)
			$L = 25$			$L = 15$		
		PWLS	62	3.5	3.0	70	0.4	2.9
		HIPOD	14	82.1	0.5	47	0.9	1.1
	15		$L = 10$			No L		
		PWLS	71	0.4	2.7	70	0.4	2.6
100		HIPOD	47	0.8	1.1	45	0.9	1.1
			$L = 25$			$L = 15$		
		PWLS	58	7.1	2.6	73	0.4	2.7
		HIPOD	4	88.8	0.5	49	0.8	1.2
	50		$L = 10$			No L		
		PWLS	72	0.4	2.4	57	0.6	2.3
		HIPOD	46	0.8	1.2	36	1.0	1.2
			$L = 25$			$L = 15$		
		PWLS	30	53.4	5.5	68	0.3	5.1
		HIPOD	0	99.5	0.5	30	18.6	1.1
	15		$L = 10$			No L		
		PWLS	72	0.2	5.1	68	0.3	5.1
150		HIPOD	38	0.9	1.3	34	0.9	1.4
			$L = 25$			$L = 15$		
		PWLS	28	61.3	5.6	76	0.2	5.7
		HIPOD	0	99.4	0.6	32	39.8	1.1
	50		$L = 10$			No L		
		PWLS	84	0.1	5.9	66	0.2	5.7
		HIPOD	50	0.6	1.6	28	0.9	1.5
			$L = 25$			$L = 15$		
		PWLS	0	53.4	5.5	30	49.7	7.6
		HIPOD	0	99.5	0.5	0	99.8	0.2
	15		$L = 10$			No L		
		PWLS	56	0.5	3.3	58	0.4	3.1
200		HIPOD	40	8.6	1.7	46	0.7	1.9
			$L = 25$			$L = 15$		
		PWLS	0	93.5	7.6	16	52.8	6.9
		HIPOD	0	99.9	0.7	0	99.9	0.3
	50		$L = 10$			No L		
		PWLS	46	0.8	2.2	24	0.7	2.2
		HIPOD	46	20.1	1.9	24	0.7	2.3

Case 1: Set $g(v) = |v|$, and this true variance function is applied in the H-PWLS.

Case 2: Set $g(v) = e^{|v|}$, but a mis-specified variance function $g(v) = \sqrt{|v|}$ is used.

We evaluate the performance of the extended PWLS for heterogeneous model (H-PWLS) by comparing it with the PWLS and the HIPOD. For the H-PWLS, we use $\widehat{\boldsymbol{\vartheta}}^{(0)} = (1, 0)$ as an initial estimate of $\boldsymbol{\theta}$ for solving (14), and we use the same way to determine initial values $\boldsymbol{\beta}^{(0)}$ and $\mathbf{w}^{(0)}$ and select tuning parameter $\widehat{\lambda}$ for solving (15) as in the PWLS. We repeat the simulation 1,000 times and use the same measurements, JD, M, and S, to summarize the results from these repetitions. These results are reported in Table 2.

From Table 2, we see that the H-PWLS significantly outperforms both the PWLS and HIPOD in all settings; it has much higher outlier detection rate measured in JD, and at the same time has much smaller masking error and swamping error measured in M and S. Especially, when $r = 20$, the H-PWLS detects almost all the outliers correctly with very small swamping error, whereas neither the PWLS nor the HIPOD works well. The H-PWLS is also illustrated in detail in Figure 3 using the data of the first repetition from the setting where $r = 10$ and $p = 15$. In addition, H-PWLS still performs consistently better than both the PWLS and HIPOD when the variance function is mis-specified. It means that the H-PWLS is relatively robust to the mis-specification of the variance function.

5.2 Real data applications

We now apply the PWLS to three real datasets: Coleman data, Salinity data, and Real Estate data. We use the random weighting method to select the tuning parameter and produce outlier probabilities for all observations.

5.2.1 Coleman Data

The Coleman data were obtained from a sample studied by Coleman et al. (1966) and further analyzed by Mosier (1997) and Rousseuw & Leroy (1987). The data include six different measurements of 20 Schools from the Mid-Atlantic and New England States. These measurements are: (1) staff salaries per pupil (*salaryP*), (2) percent of white-collar fathers (*fatherWc*), (3) socioeconomic status composite deviation (*sstatus*), (4) mean teachers verbal test score (*teacherSc*),

Table 2: Example 2 − Outlier detection evaluation for heterogeneous model (M: the mean masking probability; S: the mean swamping probability; JD: the joint outlier detection rate)

	p	Method	JD (%)	M (%)	S (%)	JD (%)	M (%)	S (%)
Case 1	15			$r = 20$			$r = 15$	
		H-PWLS	94	0.7	0.6	91	1.0	0.7
		PWLS	35	10.4	1.6	26	13.3	1.8
		HIPOD	39	9.1	2.8	36	9.7	4.9
				$r = 10$			$r = 5$	
		H-PWLS	85	1.7	0.8	20	15.2	1.5
		PWLS	13	18.9	2.3	2	34.8	4.6
		HIPOD	24	13.9	6.0	3	31.9	5.9
	50			$r = 20$			$r = 15$	
		H-PWLS	82	2.0	0.6	77	2.7	0.7
		PWLS	31	11.1	1.7	23	14.7	1.8
		HIPOD	38	9.3	3.2	33	10.2	5.2
				$r = 10$			$r = 5$	
		H-PWLS	74	3.0	0.8	12	18.6	1.6
		PWLS	13	19.2	2.3	2	37.6	4.7
		HIPOD	22	14.1	5.8	2	35.3	5.6
Case 2	15			$r = 20$			$r = 15$	
		H-PWLS	88	3.5	0.5	91	1.1	0.6
		PWLS	90	2.5	0.6	68	3.8	0.4
		HIPOD	94	0.7	2.2	84	1.8	2.2
				$r = 10$			$r = 5$	
		H-PWLS	75	2.9	0.7	14	17.3	1.5
		PWLS	41	8.2	0.7	7	25.0	2.1
		HIPOD	63	4.8	2.2	7	23.7	2.2
	50			$r = 20$			$r = 15$	
		H-PWLS	82	2.0	0.2	80	2.2	0.2
		PWLS	70	3.4	0.4	64	4.4	0.5
		HIPOD	82	2.0	2.2	72	3.2	2.2
				$r = 10$			$r = 5$	
		H-PWLS	67	4.2	0.3	13	20.9	1.1
		PWLS	43	7.8	0.7	8	28.3	2.1
		HIPOD	59	5.3	2.2	8	28.0	2.1

(5) mean mothers educational level (*motherLev*), and (6) verbal mean test score of all six graders (the outcome variable y). One wants to estimate the verbal mean test score from 5 other measurements using linear regression.

The PWLS analysis results are plotted in Figure 4, where both the weight solution path and outlier probabilities along a sequences of λ are plotted in the top and bottom panels, respectively. The vertical line is at the selected tuning parameter $\widehat{\lambda}$. This also applies for all plots Figure 4. The PWLS weights solution path (top-left panel) tells how weight $\widehat{w}_i(\lambda)$ changes with tuning parameter λ for observation i, $1 \leq i \leq 20$. From the PWLS analysis, we suggest to downweight observations 3rd, 17th, and 18th in the regression analysis.

The outlier probability plot (bottom-left panel) shows the trajectory of outlier probability $\widehat{P}_i^o(\lambda)$ for those λ. Both the 3rd and 18th observations are very likely to be outliers since the outlier probabilities are 0.99 around the vertical line. Comparing with other observations with less than 50% outlier probabilities, the 17th observation stands out with a much higher outlier probability of 0.86 around the vertical line. However, the HIPOD claims four additional outliers.

The corresponding regression estimation results are summarized at the top of Table 3 and it shows that both *sstatus* and *teacherSc* are positively associated with the outcome, while *motherLev* is negatively associated with it. Results from both the HIPOD is also listed as a comparison. The HIPOD turns to choose more outliers than the PWLS.

Table 3: Robust Estimation Results by PWLS and HIPOD for all three data sets

		Int.	*salaryP*	*fatherWc*	*sstatus*	*teacherSc*	*motherLev*
Coleman	PWLS	32.097	-1.644	0.079	0.656	1.110	-4.149
	HIPOD	1.928	-1.613	0.030	0.609	1.485	-0.489
		Int.	x_1	x_2	x_3		
Salinity	PWLS	16.913	0.711	-0.134	-0.571		
	HIPOD	0.004	1.164	-0.076	-0.049		
		Int.	Year	Area	Story	Land	
Real Estate	PWLS	1.228	0.006	0.377	0.042	0.061	
	HIPOD	1.002	0.010	0.296	0.008	0.082	

Figure 4: First column: Coleman Data; Second column: Salinity Data; Third column: Real Estate Data. Top row: PLWS weight solution paths; Bottom row: Outlier probability paths

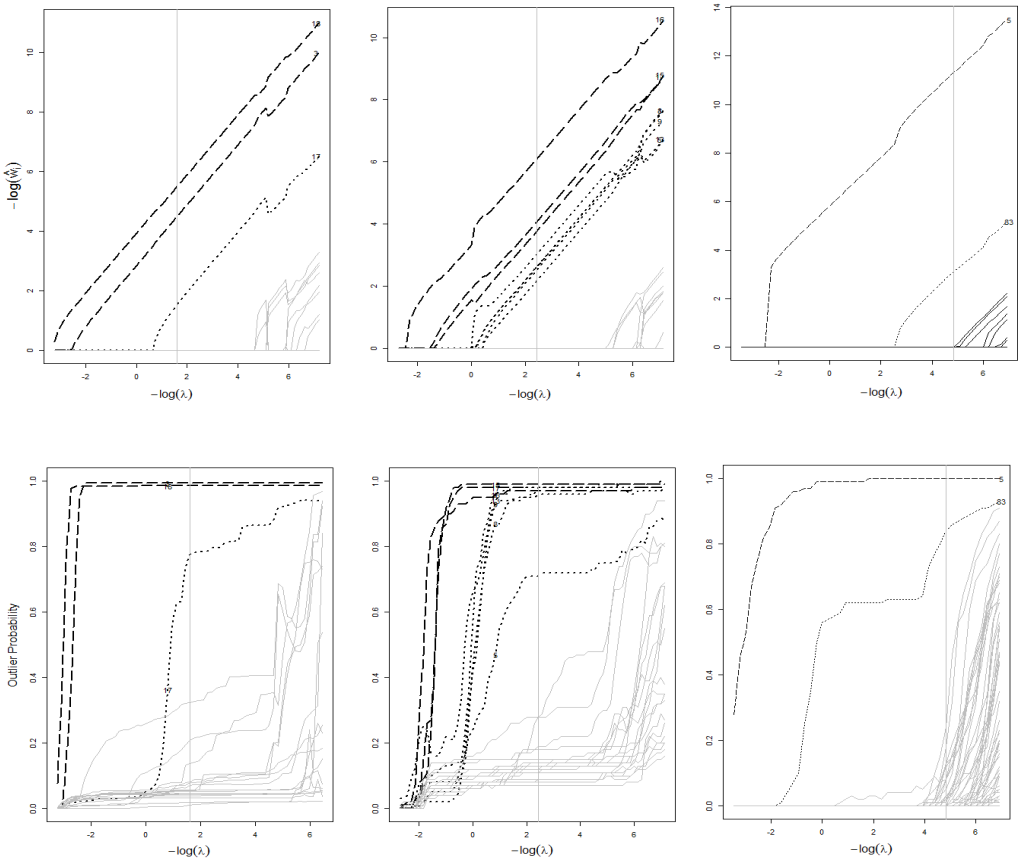

5.2.2 Salinity Data

The Salinity data consists of 28 sample points of water salinity (i.e., its salt concentration) and river discharge taken in North Carolinas Pamlico Sound. The data have four measurements: (1) lagged salinity (x_1), (2) trend (x_2), (3) discharge (x_3), and the outcome variable, salinity (y). The data were analyzed by Ruppert & Carroll (1980) and Carroll et al. (1995). Carroll et al. (1995) found that the 5th observation was masked by the 13th and 16th observations, which are corresponding to two periods of very heavy discharge.

The PWLS analysis results are reported in second column of Figure 4. The weight solution path (top-middle panel) suggests eight observations should be downweighted in the the regression analysis, the 1st, 5th, 8th, 9th, 13th, 15th, 16th, and 17th. Since there are 8 suspected outliers, some subgroup analysis on them might also be helpful. To understand these 8 suspected outliers in more detail, we examined the outlier probability plot (bottom-middle panel) carefully. All, except observation 5, among these 8 suspected outliers have very high outlier probabilities (> 0.9). The outlier probability of the 5th observation is about 0.7, is also much higher than and one for the remaining 20 observations (< 0.3).

Using the HIPOD, 15 out of 28 samples (observations 4, 6, 7, 9, 10, 11, 13, 14, 15, 17, 18, 19, 21, 23 and 24) are outliers, while observation 5 is masked.

The PWLS regression estimation results and its corresponding comparisons with HIPOD are summarized in the middle of Table 3. It shows that x_1 (lagged salinity) is positively associated with the outcome, while x_3 (discharge) is negatively associated with it.

5.2.3 Real Estate Data

The Real Estate data is taken from a Wake County, North Carolina real estate database in 2008 (Woodard & Leone 2008). This original data set includes 11 recorded variables for 100 randomly selected residential properties in the Wake County registry denoted by their real estate ID number.

The aim of the study is to predict the total assessed value of the property (logTotal: log(total value/10,000)) using four variables: the listed year in which the structure was built (Year: year built-1900), the area of the floor plan in square feet (Area: in 1000 square feet), the number of stories the structure has (Stories), and the assessed value of the land (Land, in \$10,000). Three properties (ID number 36, 50 and 58) are removed from the study since they have 0 land values.

The PWLS analysis results are reported in the third column of Figure 4. From the weight solution path (top-right panel), observations 5 and 83 (real estate ID number 86) are claimed to be outliers. Both of those two observations have considerably larger outlier probabilities (1.0 and 0.84) than the other observations (less than 0.2). See the bottom-right panel at Figure 4. The HIPOD obtains the same outlier set as the PWLS.

Robust estimation results from PWLS and HIPOD are summarized at the bottom of Table 3 and it shows both methods providing consistent results for this data set.

6 Discussion

In this manuscript, we propose a new approach to analyzing data with contamination. By assigning each observation an individual weight and incorporating an adaptive lasso-type penalty on the log-transformation of the weight vector, the aPWLS is able to perform outlier detection and robust regression.

However, like any existing penalized approach, the problem of tuning parameter selection in the aPWLS is notorious. On the one hand, the selection of tuning parameter plays an extremely important role, because it determines the number of suspected outliers. On the other hand, there is no gold standard on how to select the tuning parameter. In this manuscript, we propose two tuning methods, BIC and random weighting. The BIC was used widely in the literature and also used by She & Owen (2011) for tuning IPOD. Random weighting is a new idea for tuning parameter selection. Based on our limited numerical experience, there is not much difference in the performance between these two tuning methods, but the random weighting method can provide for each observations the probability of its being an outlier. As demonstrated using the HBK data and two real datasets, these outlier probabilities are useful for visualizing the performance of the PWLS as the tuning parameter changes.

Robust regression with variable selection has attracted much attention lately in high-dimensional data analysis. See, for example, the adaptive Lasso penalty under ℓ_1 loss in Wang et al. (2007), Huber's loss in Lambert-Lacroix & Zwald (2011) and the least trimmed squares loss in Alfons & Gelper (2013). A huge literature review on variable selection can be found in Hastie et al. 2009. Actually, we could also conduct variable selection and outlier detection simultaneously, by adding an extra penalty on the regression coefficients, say $\lambda_2 \sum_{j=1}^{p} |\beta_j|$, to the objective function

of (2).

Moreover, it is important to point out that the extended aPWLS proposed in Section 4 is actually a variation of the classical WLS aiming for fitting heterogeneity data, with the main goal being outlier detection. We consider the variance function linear models (VFLM) in Section 4, which is more general than the heterogenous model behind the classical WLS, and we can further extend the aPWLS to any variance function models.

Appendix

Proof of Theorem 1

The proof is similar to She and Owen (2011). Due to the convexity properties, we only need to check the equivalence of joint KKT functions under both (7) and (8).

We first consider joint KKT equations under (7),

$$\begin{cases} \sum_{i=1}^{n} \mathbf{x}_i' \psi((y_i - \mathbf{x}_i'\boldsymbol{\beta})/\sigma; \lambda) = 0, \\ \dfrac{\partial}{\partial \sigma} \left(\sum_{i=1}^{n} \rho((y_i - \mathbf{x}_i'\boldsymbol{\beta})/\sigma; \lambda) + 2cn \log \sigma\right) = 0. \end{cases}$$

We can check

$$\frac{\partial}{\partial \sigma}(\rho(t/\sigma, \lambda) + cn \log(\sigma)) = \begin{cases} -\lambda/\sigma + 2cn/\sigma & \text{if } |t| > \sqrt{\lambda/2}\sigma, \\ -2t^2/\sigma^3 + 2cn/\sigma & \text{if } |t| \leq \sqrt{\lambda/2}\sigma. \end{cases} \tag{18}$$

Replace t by $y_i - \mathbf{x}_i'\boldsymbol{\beta}$ and set (18) to be 0, we have

$$cn - \sum_{i \in \widehat{\mathcal{O}}} \lambda/2 - \sum_{i \in \widehat{\mathcal{O}}^c} (y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}})^2/\sigma^2 = 0,$$

where $\widehat{\mathcal{O}} = \{1 \leq i \leq n : |y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}}| > \sqrt{\lambda/2}\sigma\}$ and $\widehat{\mathcal{O}}^c = \{1 \leq i \leq n : |y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}}| \leq \sqrt{\lambda/2}\sigma\}$. Denote $r_i = y_i - \mathbf{x}_i'\boldsymbol{\beta}$ and $\mathbf{r} = (r_1, \cdots, r_n)'$. Then $\widehat{\mathbf{r}}_{\widehat{\mathcal{O}}^c}$ is the sub-vector of $\widehat{\mathbf{r}}$ for all observations in $\mathcal{O}^c$. Thus, we have

$$\widehat{\sigma}^2 = \|\widehat{\mathbf{r}}_{\widehat{\mathcal{O}}^c}\|_2^2/(cn - (\lambda/2)\sharp\{\widehat{\mathcal{O}}\}), \tag{19}$$

where $\sharp\{\widehat{\mathcal{O}}\}$ is the cardinal value of set $\widehat{\mathcal{O}}$. We now consider joint KKT of the penalized objective function in (8). From the derivative on $\mathbf{w}$,

$$2(w_i/\sigma^2)(y_i - \mathbf{x}_i'\boldsymbol{\beta})^2 + \lambda \text{sgn}(\log(w_i))(1/w_i) = 0.$$

We obtain

$$w_i = \begin{cases} \sqrt{\lambda/2}(\sigma/|\widehat{r}_i|) & \text{if } |\widehat{r}_i| > \sigma\sqrt{\lambda/2}, \\ 1 & \text{if } |\widehat{r}_i| \leq \sigma\sqrt{\lambda/2}. \end{cases} \tag{20}$$

From the derivative on σ,

$$cn\sigma^2 = \sum_{i=1}^{n} w_i^2(y_i - \mathbf{x}_i\widehat{\boldsymbol{\beta}})^2 = \sum_{i\in\widehat{\mathcal{O}}^c} \widehat{r}_i^2 + \sum_{i\in\widehat{\mathcal{O}}} (\widehat{w}_i^2\widehat{r}_i^2). \tag{21}$$

Combining with (20) and (21), we can also obtain (19).

Finally, plugging in (20) in (8), we are able to obtain the concomitant M-estimation ρ function in (7). □

References

Alfons, A., C. C. & Gelper, S. (2013), 'Sparse least trimmed squares regression for analyzing high-dimensional large data sets', *Annals of Applied Statistics* **7**(1), 226–248.

Atkinson, A., Riani, M. & Cerioli, A. (2003), *Exploring Multivariate Data with the Forward Search*, Springer.

Bach, F. (2004), Bolasso: model consistent lasso estimation through the bootstrap., *in* 'Proc. 25th Int. Conf. Machine Learning', Association for Computing Machinery, New York, pp. 33–40.

Boente, G., Pires, A. M. & Rodrigues, I. M. (2002), 'Influence functions and outlier detection under the common principal components model: A robust approach', *Biometrika* **89**(4), 861–875.

Carroll, R. J. (2003), 'Variances are not always nuisance parameters', *Biometrics* **59**(2), 211–220.

Carroll, R. J. & Härdle, W. (1989), 'Second order effects in semiparametric weighted least squares regression', *Statistics* **2**, 179–186.

Carroll, R. J., Ruppert, D. & Stefanski, L. A. (1995), 'Transformations in regression: A robust analysis', *Technometrics* **27**, 1–12.

Carroll, R. & Ruppert, D. (1988), *Transformation and Weighting in Regression*, Chapman& Hall, New York.

Cohen, J. (1960), 'A coefficient of agreement for nominal scales', *Educational and Psychological Measuremen* **20**, 37–46.

Coleman, J. S., Campbell, E. Q., Hobson, C. J., McPartland, J., Mood, A. M., Weinfeld, F. D. & York, R. L. (1966), *Equality of Educational Opportunity*, Vol. 2, Office of Education U.S. Department of Health, Washington, D.C.

Croux, C., Rousseeuw, P. & H*ö*ssjer, O. (1994), 'Generalized s-estimators', *Journal of the American Statistical Association* **89**(428), 1271–1281.

Davidian, M. & Carroll, R. J. (1987), 'Variance function estimation', *Journal of the American Statistical Association* **82**, 1079–1091.

Fang, Y. & Zhao, L. (2006), 'Approximation to the distribution of lad-estimators for censored regression by random weighting method', *Journal of Statistical Planning and Inference* **136**, 1302–1316.

Gervini, D. & Yohai, V. (2002), 'A class of robust and fully efficient regression estimators', *Annals of Statistics* **30**(2), 583–616.

Gorski, J., Pfeuffer, F. & Klamroth, K. (2007), 'Biconvex sets and optimization with biconvex functions: a survey and extensions', *Mathematical Methods of Operations Research* **66**.

Hadi, A. S., Imon, A. H. M. R. & Werner, M. (2009), 'Detection of outliers', *Wiley Interdisciplinary Reviews: Computational Statistics* **1**(1), 57–70.

Hall, P. & Carroll, R. J. (1989), 'Variance function estimation in regression - the effect of estimating the mean', *Journal of the Royal Statistical Society Series B-Methodological* **51**(1), 3–14.

Hastie, T., Tibshirani, R. & Friedman, J. (2009), *The Elements of Statistical Learning; Data mining, Inference and Prediction*, Springer Verlag, New York.

Hawkins, D., Bradu, D. & Kass, G. (1984), 'Location of seeral outliers in multiple-regression data using elemental sets', *Technometrics* **26**, 197–208.

Huber, P. J. (1973), 'Robust regression: asymptotics, conjectures and monte carlo', *Annals of Statistics* **1**(5), 799–821.

Lambert-Lacroix, S. & Zwald, L. (2011), 'Robust regression through the Huber's criterion and adaptive lasso penalty', *Electronic Journal of Statistics* **5**, 1015–1053.

Lian, H., Liang, H. & Carroll, R. J. (2014), 'Variance function partially linear single-index models', *Journal of the Royal Statistical Society, Series B* **in press**.

Ma, Y., Chiou, J.-M. & Wang, N. (2006), 'Efficient semiparametric estimator for heteroscedastic partially linear models', *Biometrika* **93**, 75–84.

Ma, Y. & Zhu, L. (2012), 'Doubly robust and efficient estimators for heteroscedastic partially linear single-index models allowing high dimensional covariates', *Journal of the Royal Statistical Society: Series B* **75**, 305–322.

Maronna, R., Martin, D. & Yohai, V. (2006), *Robust Statistics: Theory and Methods.*, Wiley.

Meinshausen, N. & Buhlmann, P. (2010), 'Stability selection (with discussion)', *Journal of the Royal Statistical Society, Series B* **72**, 417–473.

Mosier, D. (1997), 'Data analysis and regression: A second course in statistics'.

Rousseeuw, P. J. (1984), 'Least median of squares regression', *Journal of the American Statistical Association* **79**(388), 871–880.

Rousseeuw, P. & Yohai, V. (1984), *Robust regression by means of S-estimators.*, Vol. 26 of *In Robust and nonlinear time series analysis, Lecture Notes in Statistics*, Springer, New York.

Rousseuw, P. & Leroy, A. (1987), *Robust Regression and Outlier Detection*, New York: Wiley.

Rubin, D. (1981), 'The bayesian bootstrap', *Applied Statistics* **9**, 130–134.

Ruppert, D. & Carroll, R. (1980), 'Trimmed least squares estimation in the linear model', *Journal of the American Statistical Association* **5**, 828–838.

Schwarz, G. (1978), 'Estimating the dimension of a model', *The Annals of Statistics* **6**, 461–464.

Shao, J. & Tu, D. (1995), *The Jackknife and Bootstrap*, Springer Series in StatisticsI.

She, Y. & Owen, A. (2011), 'Outlier detection using nonconvex penalize regression', *Journal of the American Statistical Association* **106**(494), 626–639.

Siegel, A. F. (1982), 'Robust regression using repeated medians', *Biometrika* **69**(1), 242–244.

Sun, W., Wang, J. & Fang, Y. (2013), 'Consistent selection of tuning parameters via variable selection stability.', *Journal of Machine Learning Research* **14**, 3419–3440.

Tibshirani, R. (1996), 'Regression shrinkage and selection via the lasso', *Journal of the Royal Statistical Society, Series B* **58**, 267–288.

Wang, H., Li, G. & Jiang, G. (2007), 'Robust regression shrinkage and consistent variable selection through the lad-lasso', *Journal of Business and Economic Statistics* **25**, 347–355.

Weisberg, S. (2005), *Applied Linear Regression*, Wiley Series in Probability and Statistics, third edn, Wiley-Interscience [John Wiley & Sons], Hoboken, NJ.

Woodard, R. & Leone, J. (2008), 'A random sample of wake county, north carolina residential real estate plots', *Journal of Statistics Education* **16**(3).

Yohai, V. J. (1987), 'High breakdown-point and high efficiency robust estimates for regression', *Annals of Statistics* **15**(2), 642656.

Zou, H. (2006), 'The adaptive lasso and its oracle properties', *Journal of the American Statistical Association* **101**, 1418–1429.